Lecture 2B
MTH6102: Bayesian Statistical Methods

Eftychia Solea

Queen Mary University of London

2023

Today's agenda

Today's lecture

- Review
- Use Bayes' theorem to compute posterior pmf with discrete pmf priors.
- Use Bayes' theorem to compute posterior pfd with continuous pdf priors

Bayesian inference

- Probability model $p(y \mid \theta)$ depends on a set of parameters θ.
- θ is unknown and we would like to learn about θ.
- Let y be the observed data, assumed to be generated by this probability model, $p(y \mid \theta)$
- In Bayesian statistics, we assign probabilities on both the parameters θ and data y

Bayesian inference

- So we start with a probability distribution for the parameters $p(\theta)$, called the prior distribution.
- θ is either discrete or continuous random variable. Hence, $p(\theta)$ is either a pmf or pdf
- The prior is a subjective distribution, based on experimenter's belief, and is formulated before the data y are seen.

Bayesian inference

- Let y be the observed data.
- We then update the prior distribution (pmf or pdf) to a posterior distribution (pmf or pdf) for θ, $p(\theta \mid y)$, using Bayes' theorem

$$p(\theta \mid y) = \frac{p(\theta)\, p(y \mid \theta)}{p(y)},$$

  where the observed data enters through the likelihood $p(y \mid \theta)$.
- $p(y)$ is the normalising constant, which is given by

$$p(y) = \int p(\theta')\, p(y \mid \theta')\, d\theta' \quad \text{or} \quad \sum_{\theta'} p(\theta')\, p(y \mid \theta')$$

- $p(y)$ does not depend on θ

What does it mean?

$$p(\theta \mid y) \propto p(\theta)\, p(y \mid \theta) \qquad (1)$$

Posterior $\propto$ prior $\times$ likelihood

- $p(y \mid \theta)$ is the likelihood and it the probability of data y given the true θ.
- Start with initial beliefs/information about θ, $p(\theta)$ - this is the prior distribution formulated before the data are seen.
- Bayesian updating: Update the prior distribution using the data y, using (1).
- The updated prior, $p(\theta \mid y)$ is called the posterior distribution .
- We base our inferences about θ based on this posterior distribution.

Bayesian updating with discrete data, discrete prior

- parameter θ discrete with values θ_1 and θ_2 and prior pmf $p(\theta)$
- Discrete data x
- Discrete likelihood, $p(x|\theta)$
- Posterior pmf: $p(\theta_1|x)$, $p(\theta_2|x)$

Hypothesis	prior	likelihood	Bayes numerator	posterior
θ	$p(\theta)$	$p(x|\theta)$	$p(x|\theta)p(\theta)$	$p(\theta|x)$
θ_1	$p(\theta_1)$	$p(x|\theta_1)$	$p(x|\theta_1)\ p(\theta_1)$	$p(\theta_1|x)$
θ_2	$p(\theta_2)$	$p(x|\theta_2)$	$p(x|\theta_2)\ p(\theta_2)$	$p(\theta_2|x)$
Total	1	NOT SUM TO 1	$p(x)$	1

- Law of total probability: $p(x) = p(x|\theta_1)p(\theta_1) + p(x|\theta_2)p(\theta_2)$.
- Bayes' theorem: $p(\theta_1|x) = \frac{p(x|\theta_1)p(\theta_1)}{p(x)}$, $\quad p(\theta_2|x) = \frac{p(x|\theta_2)p(\theta_2)}{p(x)}$

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{total prob. of data}}.$$

Board Question: Coins

- There are three types of coins which have different probabilities of heads
  - Type A coins are fair, with probability 0.5 of heads.
  - Type B are bent and have probability 0.6 of heads.
  - Type C are bent and have probability 0.9 of heads.

  Suppose I have a drawer containing 5 coins: 2 of type A, 2 of type B, and 1 of type C. I pick a coin at random, and without showing you the coin I flip it once and get heads.
- Make a Bayesian update table and compute the posterior pmf that the chosen coin is each of the three coins.

- In the previous lecture, we have done Bayesian updating when we had a finite number of hypotheses or a discrete parameter θ e.g.,
  - in the diagnostic example had 2 hypotheses (HIV +ve, HIV -ve),
  - in the coin example we had 3 hypothesis (A, B and C).

- In this topic we will study Bayesian updating where there is a continuous range of hypotheses, i.e., θ is a continuous random variable.

- The Bayesian updating will be essentially the same, based on the Bayes' theorem

$$\text{posterior} \propto \text{prior} \times \text{likelihood}$$

Examples with continuous parameters

- Suppose we have a medical treatment for a disease than can succeed or fail with probability q. Then q is a continuous quantity between 0 and 1.

- The lifetime of a certain light bulb T is modeled as an exponential distribution $\exp(\lambda)$ with unknown λ. We can assume that λ takes any value greater than 0.

Baysian updating: Discrete likelihoods, continuous priors

- θ : continuous parameter with prior pdf $p(\theta)$ and range $[a, b]$.
- x : random discrete data
- discrete likelihood: $p(x|\theta)$
- posterior pdf: $p(\theta|x)$

- By **Bayes' theorem** we update the prior pdf to a posterior pdf

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)} = \frac{p(x|\theta)p(\theta)}{\int_a^b p(x|\theta)p(\theta)d\theta}.$$

- Law of total probability: $p(x) = \int_a^b p(x|\theta)p(\theta)\, d\theta$.

Baysian updating: Discrete likelihoods, continuous priors

- $p(x)$ does not depend on θ and serves as the normalising constant so that $p(\theta|x)$ is a proper pdf and integrates to 1.

- Hence, we can express Bayes' theorem in the form

$$p(\theta|x) \propto p(x|\theta)p(\theta).$$

posterior $\propto$ prior $\times$ likelihood

Bayesian inference

$$p(\theta \mid x) \propto p(\theta)\, p(x \mid \theta)$$

- $p(\theta)$ - initial beliefs/information about θ, the prior pdf.
- $p(x \mid \theta)$ - the likelihood for observed data x with parameters θ.
- Update information about θ using the likelihood.
- The resulting pdf $p(\theta \mid x)$ is called the posterior pdf of θ

$$\text{posterior} \propto \text{prior} \times \text{likelihood}$$

Bayesian updating: Discrete likelihoods, continuous priors

- Sometimes, it is better to use $p(\theta)d\theta$ to work with probabilities instead of densities
  e.g the prior probability that θ is in a small interval of width $d\theta$ around 0.5 if $p(0.5)d\theta$.

- In this case, the Bayes' theorem is

$$p(\theta|x)d\theta = \frac{p(x|\theta)p(\theta)d\theta}{p(x)} = \frac{p(x|\theta)p(\theta)d\theta}{\int_a^b p(x|\theta)p(\theta)\,d\theta}.$$

Bayesian updating: Discrete likelihoods, continuous priors

- θ : continuous parameter with prior pdf $p(\theta)$ and range $[a, b]$.
- x : random discrete data
- likelihood: $p(x|\theta)$

Bayesian updating table

Hypothesis	prior prob	likelihood	Bayes numerator	posterior prob. $p(\theta|x)d\theta$
θ	$p(\theta)d\theta$	$p(x|\theta)$	$p(x|\theta)p(\theta)d\theta$	$\frac{p(x|\theta)p(\theta)d\theta}{p(x)}$
Total	$\int_a^b p(\theta)d\theta = 1$		$p(x) = \int_a^b p(x|\theta)p(\theta)d\theta$	1

- The posterior density $p(\theta|x)$ is obtained by removing $d\theta$ from the posterior probability in the table.

Example: Binomial data, Beta prior

- A biased coin has probability of heads q which is unknown.
- We toss the coin n times and observe k heads (This is my data $x = k$).
- The binomial likelihood for this problem:

$$p(k \mid q) = \binom{n}{k} q^k (1-q)^{n-k}$$

- For Bayesian inference, we need to specify a prior distribution for q.
- q is a continuous quantity between 0 and 1.
- What is a possible probability distribution for q (or family of distributions)?

Example: Binomial data, Beta prior

- The family of Beta distributions seems a natural choice for a prior distribution for q, since it describes continuous random variables with support on $[0, 1]$.
- If $q \sim \text{Beta}(\alpha, \beta)$, its probability density function is

$$f(q) = \frac{q^{\alpha-1}(1-q)^{\beta-1}}{B(\alpha,\beta)},\ 0 \leq q \leq 1,$$

where B is the Beta function and α and β are parameters,

$$B(\alpha,\beta) = \int_0^1 x^{\alpha-1}(1-x)^{\beta-1}\, dx$$

$$B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$$

Beta distributions

- Probability density functions.
- If $q \sim \text{Beta}(\alpha, \beta)$

$$E(q) = \frac{\alpha}{\alpha + \beta}$$

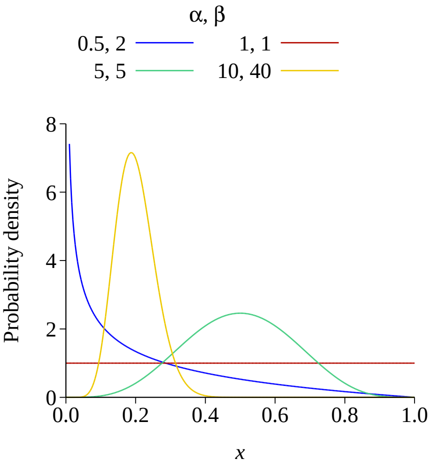

Example: Binomial data, Beta prior

- Bayesian updating:

$$\text{posterior} \propto \text{prior} \times \text{likelihood}$$

- The posterior distribution $p(q|x)$ is proportional to

$$p(q \mid k) \propto q^{k+\alpha-1}(1-q)^{n-k+\beta-1}$$

- We recognise this to have the form of a beta distribution, so the posterior is a beta distribution, $\text{beta}(k+\alpha, n-k+\beta)$.

- Hence, the normalising constant must be $1/B(k+\alpha, n-k+\beta)$.

Bayesian updating: Discrete likelihoods, continuous priors

- The actual, normalized pdf is

$$p(q \mid k) = \frac{q^{k+\alpha-1}(1-q)^{n-k+\beta-1}}{B(k+\alpha, n-k+\beta)},$$

the pdf of a $\text{Beta}(k+\alpha, n-k+\beta)$ r.v. (**Remember:** the random variable is q and k is fixed).

- Bayesian updating: We update the prior $\text{Beta}(\alpha, \beta)$ to posterior $\text{Beta}(k+\alpha, n-k+\beta)$.

Bayesian updating table: Discrete likelihoods, continuous priors

Example: Binomial data, Beta prior

- $Y \sim$ Binom(n,q), with q unknown
- Continuous hypotheses q in $[0, 1]$.
- Data y
- Prior $p(q)$
- Likelihood $p(y|q)$

Hypothesis	prior prob.	likelihood	Bayes numerator	posterior prob.
q	Beta$(\alpha, \beta)dq$	binomial(n, q)	$cq^{k+\alpha-1}(1-q)^{n-k+\beta-1}dq$	Beta$(k+\alpha, n-k+\beta)dq$
Total	1		$T = \int_0^1 cq^{k+\alpha-1}(1-q)^{n-k+\beta-1}dq$	1

- The posterior density is Beta$(k+\alpha, n-k+\beta)$
- **Note:** We don't need to compute T. Once we know the posterior is of the form $cq^{k+\alpha-1}(1-q)^{n-k+\beta-1}$ we have to find c that makes it a proper density. In this case $c = 1/\text{Beta}(k+\alpha, n-k+\beta)$

Unknown parameters and prior parameters

Remarks:

- We need to distinguish between the parameters we are estimating, which we generally have denoted by θ and the parameters for the prior distribution(s).
- In this binomial example, q is uncertain: we have prior and posterior distributions for q.
- The parameters of the prior distribution, here α and β, are taken as fixed.

Board question: bent coin

- Bent coin with unknown probability θ of heads
- Prior: $p(\theta) = 2\theta$ on $[0, 1]$
- Data: toss and get heads

- Compute the Bayesian update table.
- Find the posterior pdf to this data.

Posterior mean

In Bayesian how would you choose a particular value of q?

- A natural estimate for q is the mean of the posterior distribution $p(q|k)$, called the posterior mean.
- For the binomial case with Beta(α, β) prior, the posterior mean is

$$\hat{q}_{\mathrm{B}} = E(q \mid k) = \frac{k + \alpha}{n + \alpha + \beta}.$$

- The prior distribution has mean $\alpha/(\alpha + \beta)$ which would be our best estimate of q without having observed the data.
- Ignoring the prior, we would estimate q using the maximum likelihood estimate (MLE)

$$\hat{q} = \frac{k}{n}$$

- The Bayes' estimate $\hat{q}_{\mathrm{B}}$ combines all of this information.

Posterior mean

- Note that we can rewrite $\hat{q}_{\text{b}}$ as

$$\hat{q}_{\text{B}} = \frac{n}{n+\alpha+\beta}\left(\frac{k}{n}\right) + \frac{\alpha+\beta}{n+\alpha+\beta}\left(\frac{\alpha}{\alpha+\beta}\right).$$

- Thus $\hat{q}_{\text{B}}$ is a linear combination of the prior mean and the MLE, with the weights being determined by n, α and β

Flat priors

- One important prior is called flat prior or uniform prior.
- A flat prior assumes that every hypothesis is equally probable.
- For example if q has range $[0, 1]$, then $p(q) = 1$ is a flat prior.
- E.g. a uniform distribution on $[0, 1]$ is $\text{Beta}(1, 1)$
- So, posterior distribution is $\text{Beta}(k + 1, n - k + 1)$
- Posterior mean: $E(q \mid k) = \dfrac{k + 1}{n + 2}$

Board question

- Bent coin with unknown probability θ of heads
- Flat prior: $p(\theta) = 1$ on $[0, 1]$
- Data: toss 27 times and get 15 heads and 12 tails.

- Compute the Bayesian update table.
- Give the integral for the normalising factor but do not compute it. Call its value T and give the posterior pdf in terms of T.

Board question

- A medical treatment has unknown probability θ of success.
- We assume treatment has prior $f(\theta) \sim \text{Beta}(5, 5)$.

1. Suppose you test it on 10 patients and have 6 successes. Find the posterior distribution on θ. Identify the type of the posterior pdf
2. Suppose you recorded the order of the results and got SSSFFSSSFF. Find the posterior based on this new data.